



CA Fast Unload & CA Fast Load for DB2 for z/OS

Best practices

Manuel Gómez Burriel – May, 31st, 2018 – 2.13



For Informational Purposes Only

This presentation was based on current information and resource allocations as of **May 2018** and is subject to change or withdrawal by CA at any time without notice. Notwithstanding anything in this presentation to the contrary, this presentation shall not serve to (i) affect the rights and/or obligations of CA or its licensees under any existing or future written license agreement or services agreement relating to any CA software product; or (ii) amend any product documentation or specifications for any CA software product. The development, release and timing of any features or functionality described in this presentation remain at CA's sole discretion. Notwithstanding anything in this presentation to the contrary, upon the general availability of any future CA product release referenced in this presentation, CA will make such release available (i) for sale to new licensees of such product; and (ii) to existing licensees of such product on a when and if-available basis as part of CA maintenance and support, and in the form of a regularly scheduled major product release. Such releases may be made available to current licensees of such product who are current subscribers to CA maintenance and support on a when and if-available basis. In the event of a conflict between the terms of this paragraph and any other information contained in this presentation, the terms of this paragraph shall govern.

Certain information in this presentation may outline CA's general product direction. All information in this presentation is for your informational purposes only and may not be incorporated into any contract. CA assumes no responsibility for the accuracy or completeness of the information. To the extent permitted by applicable law, CA provides this presentation "as is" without warranty of any kind, including without limitation, any implied warranties or merchantability, fitness for a particular purpose, or non-infringement. In no event will CA be liable for any loss or damage, direct or indirect, from the use of this document, including, without limitation, lost profits, lost investment, business interruption, goodwill, or lost data, even if CA is expressly advised in advance of the possibility of such damages. CA confidential and proprietary. No unauthorized copying or distribution permitted.

CA Fast Unload & Fast Load for DB2 for z/OS best practices

- Agenda

- CA Fast Unload for DB2 best practices
 - product customization
 - maintaining data integrity,
 - maximizing performance
 - reducing CPU time.
 - automating the creation of test tables
 - restarting jobs
 - Usability
- CA Fast Load for DB2 best practices
 - Sort
 - Performance
 - Usability
 - Functionality

CA Fast Unload for DB2 for z/OS best practices

- Customization: create a **PFU** parmlib member

- specify site-specific default values for CA Fast Unload® for DB2 for z/OS keywords
- PFU parmlib member can contain any CA Fast Unload® for DB2 for z/OS syntax.
- Syntax in the “**PFU**” member is the same as in job.
- Keyword values specified will work as default settings.
- Recommended
 - SORTNUM, SORTDEV
 - SQL-ACCESS EXTENSION
 - PRINT-OPTIONS PTIMSG or ALLMSGs
 - EMPTY-RC: allow RC 04 when no rows for a SELECT
 - DISPLAY-STATUS

CA Fast Unload for DB2 for z/OS best practices

- Customization: compatibility

DB2 Compatibility Matrix

Compatibility/Release	18.0	19.0	20.0
Db2 V9	NO	NO	NO
Db2 V10	YES PTFs	YES PTFs	YES PTFs
Db2 V11	YES PTFs	YES PTFs	YES PTFs
Db2 V12	NO	YES PTFs Info	YES PTFs Info

Pervasive Encryption Compatibility Matrix

Compatibility/Release	18.0	19.0	20.0
Pervasive Encryption z/OS V2R2	NO	TBD	TBD
Pervasive Encryption z/OS V2R3	NO	TBD	TBD

BM zEnterprise Compatibility Matrix

Compatibility/Release	18.0	19.0	20.0
zEC12	YES	YES	YES
zBC12	YES	YES	YES
z13/z13s	YES	YES	YES
z14	YES PTFs	YES PTFs	YES PTFs

CA Fast Unload for DB2 for z/OS best practices

- Data integrity: Unload access type

- SQL-ACCESS [**NONE**|ONLY|EXTENSION]
 - **NONE**: SELECT statement processed natively (outside DB2)
 - Best performance
 - Less impact on DB2 workload
 - Required to unload from an image copy
 - SHRLEVEL keyword specifies the data integrity level
 - ONLY: SELECT statement processed by DB2
 - EXTENSION: SELECT statement processed natively when possible, switching to DB2 when not supported.
 - Mix of NONE and ONLY options.

CA Fast Unload for DB2 for z/OS best practices

- Data Integrity: SHRLEVEL

- SHRLEVEL only applicable to SQLs executed natively.
- With data integrity
 - SHRLEVEL REFERENCE:
 - SQL-ACCESS NONE or EXTENSION.
 - Table data buffers quiesced.
 - MULTI-LOCKING: common point of consistency for unloaded data across all SELECT statements.
 - PART-INDEPENDENCE YES
 - Consider using alternative input dataset formats INPUT-FORMAT
 - IMAGECOPY,
 - DSN1COPY
 - CONCURRENT COPY
- Without data integrity: “dirty read”. Concurrent access.
 - SHRLEVEL CHANGE
 - CHANGE BP-LOOKUP, QUIESCE|NO-QUIESCE
 - SHRLEVEL IGNORE

CA Fast Unload for DB2 for z/OS best practices

- Data Integrity: considerations

- STOPPED-TS [**NO**|YES]
 - specifies whether to unload a table from a stopped tablespace.
 - Full integrity...if STOP status is under control
- REFERENCE-TS-COPY [**NO**|YES]
 - used with SHRLEVEL REFERENCE
 - unloads tablespace that is in RW, COPY status.
 - Status prevents a quiesce from being performed, therefore data pages might still be in DB2 buffers.
 - Data integrity at risk: “dirty read”.

CA Fast Unload for DB2 for z/OS best practices

- Maximize performance

- Allocate sort work datasets dynamically
 - Omit STxxWKnn and STxxMSG ddnames
 - Include SORT dynalloc parameters
 - SORTDEV device-type
 - SORTNUM # temporary datasets
 - Include ESTIMATED-ROWS in every SELECT under SELECT-OPTIONS parameter to facilitate an optimal temporary and output datasets allocation.
- Optimize input/output processing
 - COPY-BUFFERS: for IMAGCOPY and DSN1COPY zOS input datasets
 - VSAM-BUFFERS: for DB2 ESDS dataset...if SQL-ACCESS NONE.
 - IO-BUFFERS: for SYSREC output datasets
 - ...but avoid system paging

CA Fast Unload for DB2 for z/OS best practices

- Maximize performance

- Optimize SELECT statements processing:
 - Include several SELECTs in a single step
 - SQL-ACCESS ONLY / EXTENSION
 - CURRENT-DEGREE [NONE|ANY|1]
- Setting the data INPUT-FORMAT
 - TABLE:
 - required to unload LOB data
 - zIIP processing when processed natively (SQL-ACCESS NONE)
 - Temporal tables support
 - Copies: DSN1COPY / CONCURRENTCOPY (DFSMS) / FLASHCOPY / IMAGCOPY / INCREMENTAL / INLINE (Reorg)
 - DDLDDN—Retrieve DDL from a Data Set...if not consistent with catalog

CA Fast Unload for DB2 for z/OS best practices

- Reducing CPU

- Setting the data OUTPUT-FORMAT
 - Avoid data conversion parameters: DATE, TIME, fieldprocs, etc
 - LOAD: IBM Load compatibility. Best performance.
 - NEWOBID: Specify the OBID of the table to be loaded.
 - required to unload LOB data
 - zIIP processing when processed natively (SQL-ACCESS NONE)
 - Temporal tables support
 - VARIABLE: if variable length columns present. Required for XML columns
 - DSNTIAUL: IBM DSNTIAUL compatibility. FB.
 - FIXED: FB dataset, If variable length columns present, padded.
 - COMMA-DELIMITED: IBM IND\$FILE file-transfer compatibility.
 - VARIABLE,[DB2LOAD|FASTLOAD|XMLDB2V9]|EXTERNAL}
 - DSNTIAUL:
 - Avoid data conversion parameters: DATE, TIME, fieldprocs, etc

CA Fast Unload for DB2 for z/OS best practices

- CPU time saving

- INPUT FORMAT
 - TABLE
 - SQLACCESS NONE / EXTENSION
- OUTPUT-FORMAT
 - DSNTIAUL: avoid
 - INTO
 - TRIM
 - DATE/TIME types column with a LOCAL date/time exits
 - Field procedures on any column
 - ASCII, CCSID, EBCDIC, and UNICODE keywords
 - COMMA-DELIMITED / LOAD
 - DATE/TIME types column with a LOCAL date/time exits
 - Field procedures on any column
 - ASCII, CCSID, EBCDIC, and UNICODE keywords

CA Fast Unload for DB2 for z/OS best practices

- Reducing CPU: zIIP processing

- Requirements:
 - ZIIP {**YES**|NO}
 - A single SELECT statement
 - multiple SELECTs are not allowed
 - INPUT-FORMAT TABLE
 - SQL-ACCESS NONE (or EXTENSION if statement processed natively)

CA Fast Unload for DB2 for z/OS best practices

- Automate creation of tables for TEST environments

- Creation of single tables in the same or different DB2s.
- Created tables can contain some or all columns from the original table.
- These test tables can also have added column data initialized based on various criteria.
- **SELECT** statement specification
 - literals
 - columns to unload.
 - field specifications in the **INTO** clause to specify **alternate column formatting** for the unloaded data.
 - **mask** sensitive client data by including one or more **DEFAULTIF** (*condition*) **INITIAL** (*value*) clauses for each column.
 - The **LIMIT** keyword lets you limit the total number of rows to unload (sampling).

CA Fast Unload for DB2 for z/OS best practices

- Automate creation of tables for TEST environments

- **SAMPLE**: unloads a sample of data from your table.
 - **SAMPLE** {x|(,y)|(x,y)} unloads **x** rows every **y** selected rows
 - Specify at global option or within **SELECT-OPTIONS**
- **SELECT** statement specification
 - literals
 - columns to unload.
 - field specifications in the **INTO** clause to specify **alternate column formatting** for the unloaded data.
 - **mask** sensitive client data by including one or more **DEFAULTIF** (*condition*) **INITIAL** (*value*) clauses for each column.
 - The **LIMIT** keyword lets you limit the total number of rows to unload (sampling).

CA Fast Unload for DB2 for z/OS best practices

- Automate creation of tables for TEST environments

- DDL-CONTROL keyword:
 - generates CREATE TABLE DDL statements
 - Only compatible with SQL-ACCESS NONE
 - DDL-CONTROL [**NONE**|INTABLE[,ONLY]|OUTTABLE[,ONLY]|BOTH[,ONLY]]
 - Output to CTLDDN template-name | ddname
- LOAD-CONTROL keyword:
 - IBM format
 - CA Fast Load format

CA Fast Unload for DB2 for z/OS best practices

- restarting jobs

- Restart possibilities if a Fast Unload job fails:
 - At job level
 - At utility statement level
 - At utility phase level
- Execution requirements for job restarting
 - specifying AUTO-RESTART YES in *hlq.CDBAPARM(UTIL)* and
 - specifying a RESTART parameter in the EXEC statement of your failed job
 - RESTART(PHASE):
 - Restart from failed utility statement
 - Execute subsequent utility statements without a restart parameter value.
 - RESTART(BYPASS),
 - RESTART from failed utility statement
 - BYPASS; subsequent statements are skipped the BYPASS option
 - RESTART(TERM):
 - RESTART from failed utility statement
 - TERM; delete restart rows. Does not execute subsequent statements.

CA Fast Unload for DB2 for z/OS best practices

- LOB columns

- Performance and support recommendations
 - INPUT-FORMAT TABLE: required
 - Other INPUT-FORMAT options (DSN1COPY, IC,..) do not support LOB columns
 - SQL-ACCESS NONE: PFU direct reading
 - zIIP processing
 - Multitasking
 - Decrease ssidDBM1 overhead...subsystem workload
 - Possibility to specify partitions of a LOB tablespace
 - ONLY and EXTENSION also supported...in case of code character translations
 - OUTPUT-FORMAT VARIABLE
 - Compatible with IBM LOAD and CA Fast Load
 - Disk space used more effectively
 - If output dataset length > 32K: RECFM=VBS
 - OUTPUT-FORMAT EXTERNAL: no LOB support

CA Fast Unload for DB2 for z/OS best practices

- Elapsed time saving

- Multitasking
 - At partition level
 - PART
 - Partitions to unload: all, range, list....any combination
 - PART-INDEPENDENCE [NO|YES]
 - Access to non-eligible partitions to unload by other processes.
 - Skip the “global” lock at partitioned tablespace level. (native mode)
 - PART-SEPARATE {NO|YES}
 - One unload subtask per partition
 - Even at SELECT level

CA Fast Unload for DB2 for z/OS best practices

- Elapsed time saving

- Multitasking capability
 - Automatic under following syntax conditions
 - SQL-ACCESS NONE
 - INPUT-FORMAT TABLE
 - TAPE-STACK NO (or omitted)
 - Multitasking level
 - 1 unload task per tablespace: for non-partitioned tablespaces
 - 1 unload task per SELECT: for partitioned tablespaces

CA Fast Unload for DB2 for z/OS best practices

- Elapsed time saving

- Multitasking: Process Summary Report capability

```
***PSR*****
*
*           F A S T   U N L O A D
*           P R O C E S S   S U M M A R Y   R E P O R T
* QUICK FIND KEYS:
*           PSR   - PROCESS SUMMARY REPORT
*
*****
```

```
PTFU261 - TASK #01 -- MESSAGE DDNAME: PTMSG01 -- MAX RC: 00 INCLUDES:
PTFU262 - SELECT STMT#01 -- UNLDDN SYSREC01: .....800
PTFU262 - SELECT STMT#02 -- UNLDDN SYSREC02: .....800
PTFU263 - TASK #01 -- ELAPSED DISPATCH TIME: 00:00:10.51
```

```
PTFU261 - TASK #02 -- MESSAGE DDNAME: PTMSG02 -- MAX RC: 04 INCLUDES:
PTFU262 - SELECT STMT#03 -- UNLDDN SYSREC03: ..... 0
PTFU262 - SELECT STMT#04 -- UNLDDN SYSREC04: .....800
PTFU263 - TASK #02 -- ELAPSED DISPATCH TIME: 00:00:07.43
```

```
PTFU290 - TOTAL UNITS OF WORK PROCESSED: 02 - MAX CONCURRENT TASKS: 02
```

CA Fast Unload for DB2 for z/OS best practices

- Usability

- TEMPLATE: Dynamic Allocation Parameters for Output Data Sets

- Not a FASTUNLOAD keyword
- Define naming conventions for SYSREC, SYSDDL, SYSCTL
- If present in SYSIN,
 - specify before FASTUNLOAD control statement
 - Omit SYSTEMPL DD statement
- TEMPLATE *template-name*

DSN *name-expression* (accept symbolics: &DB../TS..p&PA)

[*common-options*]

[*disk-options* | *tape-options*]

CA Fast Unload for DB2 for z/OS best practices

- Usability

- **OUTPUT-FORMAT** {**LOAD**|**FIXED**|**DSNTIAUL**|**COMMA-DELIMITED**|**VARIABLE**,[**DB2LOAD**|**FASTLOAD**|**XMLDB2V9**]|**EXTERNAL**}
- **LOAD**: IBM unload format. Compatible with IBM LOAD and CA FastLoad
- **FIXED**: output fields padded to maximum length
- **DSNTIAUL** IBM DSNTIAUL format. output fields padded to maximum length
- **COMMA-DELIMITED** PC format. Delimiter character.
- **VARIABLE** VB/VBS
 - **DB2LOAD**
 - **FASTLOAD**
 - **XMLDB2V9** obsolete
- **EXTERNAL**

CA Fast Load for DB2 for z/OS best practices

- Agenda

- Sort
- Performance
- Usability
- Functionality

CA Fast Unload for DB2 for z/OS best practices

- Customization: create a PFL parmlib member

- specify site-specific default values for CA Fast Load® for DB2 for z/OS keywords
- PFL parmlib member can contain any CA Fast Load® for DB2 for z/OS syntax.
- The keyword syntax in PFL member same as in a job.
- Keyword values specified will work as default settings.
- Recommended
 - SORTNUM, SORTDEV
 - SQL-ACCESS EXTENSION
 - PRINT-OPTIONS PTIMSG or ALLMSGs
 - EMPTY-RC: allow RC 04 when no rows for a SELECT
 - DISPLAY-STATUS

CA Fast Load for DB2 for z/OS best practices

- Sorting

- Dynamic allocation
 - SORTDEV
 - SORTNUM
 - MAXTASK: [n/AUTO] concurrent sort tasks
 - The calculation is based on how many indexes the target table has
- Volume of information
 - Existing rows in target table: TOTALROWS value in SYSIBM.SYSTABLESPACESTATS (RTS) table is used. If 0 or null, the CARDF value from DB2 catalog. ESTIMATED-ROWS keyword omitted.
 - Input records to load in SYSULD: If information is compress and/or with LOB/VARCHAR data, specify ESTIMATED-INPUT keyword

CA Fast Load for DB2 for z/OS best practices

- Sorting: CA Utility Sort

- CA partners with Syncsort
 - Use Syncsort as a sort program in PRR, PFL, PLA
 - CA Utility Sort
 - API to communicate with sort program
- Advantages
 - Optimize the sort dynamically based on the workload to be sorted
 - Utilization of block mode user exits provides additional performance improvement and zIIP offload

CA Fast Load for DB2 for z/OS best practices

- Performance: Integration with other utilities

- Integration: COPY & STATS when LOAD REPLACE
 - Copy processing
 - COPY-BUFFERS: I/O buffers to allocate for CA Quick Copy.
 - COPY-EXCP:
 - Runstats processing:
 - DB2 catalog: UPDATE-CATSTATS
 - Call IBM RUNSTATS:
 - RSD: RUNSTATS TABLESPACE tsname TABLE ALL INDEX ALL
 - RSP: RUNSTATS TABLESPACE tsname USE PROFILE TABLE ALL
 - Test option [RSDT, RSPT]: REPORT YES UPDATE NONE
 - Real Time Stats: UPDTE-RTS-TABLES YES in hlq.CDBAPARM (UTIL).
 - Keeping RTS updated, application queries are optimized without additional RUNSTATS utility being executed.
 - PDA statistics: UPDATE-PDASTATS

CA Fast Load for DB2 for z/OS best practices

- Performance: Input datasets

- INPUT-FORMAT:
 - SEQ: default. FB, VB, VBS.
 - IBM DSNTIAUL unload format compatibility
 - UNLOAD
 - An IBM REORG with the UNLOAD ONLY option
 - A CA FAST UNLOAD with the OUTPUT-FORMAT LOAD option
 - A discard data set from CA Rapid Reorg® for DB2 for z/OS
 - An exceptions data set from CA Fast Check® for DB2 for z/OS
 - The column definitions of the source and target tables must match.
 - No data or character conversions occur. CPU savings.
 - INTERNAL: IBM UNLOAD job with FORMAT INTERNAL
 - DELIMITED: delimiter character for columns

CA Fast Load for DB2 for z/OS best practices

- Performance: output Db2 spaces

- **OUTPUT-CONTROL BUILD:** load TS directly from input.
 - Writes the input data from one or more SYSULD data sets directly into the tablespaces and indexspaces.
 - No intermediate data sets are used.
 - This load type reduces I/O, reduces disk storage needs and speeds processing.
 - Recommended with
 - REPLACE keyword
 - No duplicated unique index entries: discard processing avoidance.
 - No out of sequence index entries: discard processing avoidance, if RECLUSTER NO (input data expected in clustered order)
 - RECLUSTER NO with OUTPUT-CONTROL BUILD also minimizes the workspace DASD by eliminating the need for sort work data sets.

CA Fast Load for DB2 for z/OS best practices

- Performance: output Db2 spaces

- Index processing
 - IXBUFFER-SIZE: Buffer Size for Index Processing.
 - NO-DUPLCHECK—Bypass Check for Duplicate Index Keys when OUTPUT-CONTROL REORG
- Partitioned tablespaces
 - Load from different input datasets: INDDN ddname
 - Concatenate datasets under 1 ddname in case the information for a specific partition spreads into more than one dataset.
 - If LRECLs different, greater first in the concatenation

CA Fast Load for DB2 for z/OS best practices

- Performance: Unix pipes

- UNIX Pipes in zOS/USS

- Alternative to

- BMC: MainView Batch Optimizer - Job Optimizer Pipes
 - IBM BatchPipes

- BatchPipes maintains a *short* queue of records being passed between “writer” and “reader” processes, and concurrently.

- Provide better performance and cost reduction

- Elapsed and CPU time savings.
 - zIIP offload

CA Fast Load for DB2 for z/OS best practices

- Usability

- Column order specification for an INPUT-FORMAT DELIMITED process.
- process the output data from IBM UNLOAD jobs that specify FORMAT INTERNAL.
- Index building with existing duplicated keys: OUTPUT-CONTROL ALL
- Data sharing support: subgroup

CA Fast Load for DB2 for z/OS best practices

- Usability:

- File Clean-up: DELETE-FILES automatic deletion
 - only if SYSRECxx, SYSUT1xx, and SYSERR allocated manually
 - Only if LOAD completes successfully
 - Specified as a global option, where it applies to the entire job.
 - Format: DELETE-FILES [YES|NO]
 - DELETE-FILES YES: forces deletion of output datasets
 - SYSRECxx
 - SYSUT1xx
 - SYSERR: only if corresponding SYSDISC is empty
 - DELETE-FILES NO: honours value of DISP DD parameter for every dataset

CA Fast Load for DB2 for z/OS best practices

- Functionality

- Fast Load Resume Critical
 - Provide Load Resume processing without having to take the object offline.
 - Eliminate object outage when loading additional data.
- Fast Load Replace with Fast Switch
 - Eliminate object outages when reloading tables with new data.
 - Eliminate outages when replacing read only tables with refreshed data.

- Functionality: Online Load Resume...Why?

- 24/7 environment
 - Active tables
 - Data updates needed
 - High data availability required
-
- How to add information
with no business impact?

[illegible]

CA Fast Load for DB2 for z/OS best practices

- Functionality : Online Load Resume...Solution

```
//SYSIN      DD *
```

```
FASTLOAD
```

```
DATA-AVAILABLE CRITICAL
```

```
COMMIT-FREQUENCY (50000, 100)
```

```
RESUME YES
```

Triggers online load and keeps table in read/write mode during load process

Specifies how frequently to commit the loaded data, and how many rows to load at a time

Loads new data at end of table

CA Fast Load for DB2 for z/OS best practices

- Functionality: Online Load Resume...Sample JCL

```
.....
//SYSULD DD DISP=status,
//          DSN=userid.table.SYSREC
//*
//SYSIN DD *
FASTLOAD
INPUT-FORMAT UNLOAD
DATA-AVAILABLE CRITICAL
COMMIT-FREQUENCY (50000, 100)
INDDN SYSULD
INTO TABLE userid.table
RESUME YES
```

Identifies the input file

Indicates that the input file was generated by CA Fast Unload

Specifies online load processing

Specifies the commit frequency and number of rows per insert

Identifies the input data set

Specifies the table (or partitions) to update

Adds new data at the end

CA Fast Load for DB2 for z/OS best practices

- Functionality: Online Load Resume...results

- **Successful online load**

```
***** ***** Top of Data *****
PFL0805I - Table owner.table_name
PFL0810I - Records Processed .....nn
PUT0019I - Elapsed Time: hh:mm:ss - dd mmm yyyy hh:mm:ss
***** ***** Bottom of Data *****
```

- **Failed online load**

```
***** ***** Top of Data *****
PFL0206E - FASTLOAD application did not complete successfully
PFL0207E - T_Space Task ECB: ech_code Return Code: return_code
***** ***** Bottom of Data *****
```

CA Fast Load for DB2 for z/OS

- Functionality: Online Load Resume...Summary

- **Use case**
 - 24/7 environment
 - Active tables
 - Data updates needed
 - High data availability required
- **CA Fast Load for DB2**
 - DATA-AVAILABLE CRITICAL
 - COMMIT-FREQUENCY
 - RESUME YES

CA Fast Load for DB2 for z/OS

- Functionality: Online Load Replace...Summary

- Online Load Replace
 - 24/7 environment
 - Active tables initialization
 - Data refresh needed at regular intervals
 - High data availability required
- How to **replace** information
with no business impact?

Reference tables

The diagram illustrates the concept of online load replace by showing four overlapping tables. Each table has a blue header and a grid of data rows. The tables are arranged in a staggered, overlapping fashion, suggesting a sequence of updates or replacements. The top-left table is the most visible, while the others are partially obscured behind it, creating a sense of depth and continuity.

CA Fast Load for DB2 for z/OS

- Functionality: Online Load Replace...Summary

- Online Load Replace: Parametrization
 - [DATA-AVAILABLE HIGH](#) to invoke the fast switch function, create shadow copies, and keep the original tablespace in read-only (RO) mode during the load.
 - [OUTPUT-CONTROL BUILD](#) to load the data directly from the input file.
 - [RESUME NO REPLACE](#) to indicate that you are replacing the original table data (not adding more data to it).

CA Fast Load for DB2 for z/OS

- Functionality: Online Load Replace...Syntax

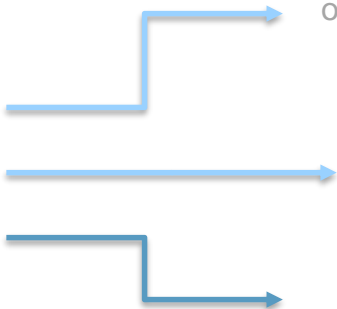
```
//SYSIN DD *
```

```
FASTLOAD
```

```
DATA-AVAILABLE HIGH
```

```
OUTPUT-CONTROL BUILD
```

```
RESUME NO REPLACE
```



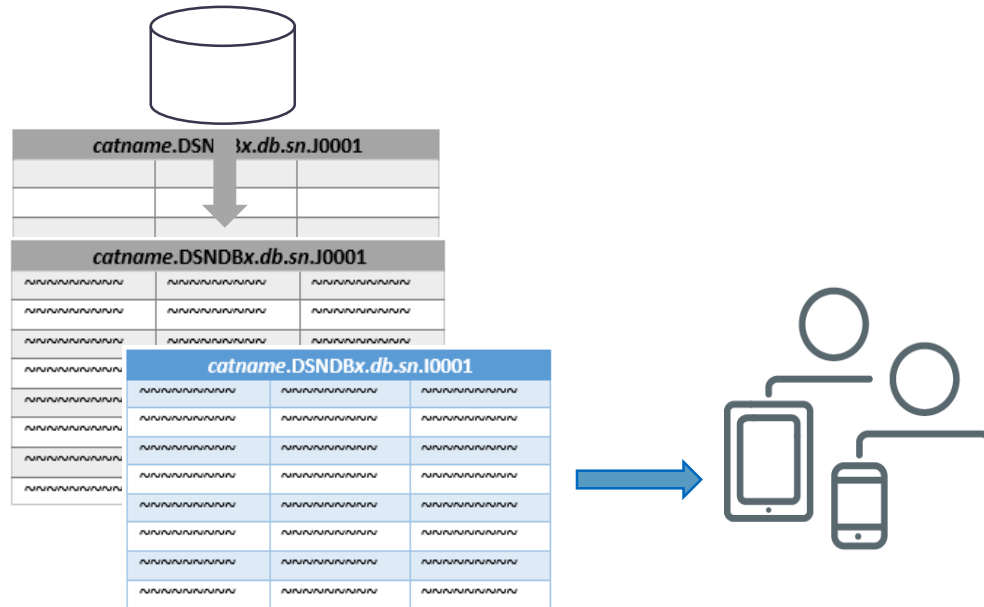
Triggers fast switch and keeps table in read-only mode during load process

Loads data directly into table

Replaces the original table data

CA Fast Load for DB2 for z/OS

- Functionality: Online Load Replace...Process



CA Fast Load for DB2 for z/OS

- Functionality: Online Load Replace...FastSwitch

- Successful fast switch: PFL0173I

Original table

A	B	C
xxx	xxx	xxx
yyy	yyy	yyy
zzz	zzz	zzz

(catname.DSNDBx.db.sn.I0001)



Shadow table

A	B	C
xxx	aaa	xxx
yyy	yyy	bbb
zzz	zzz	zzz

(catname.DSNDBx.db.sn.J0001)

CA Fast Load for DB2 for z/OS

- Functionality: Online Load Replace...Process

- Successful fast switch: PFL0173I

```
PFLPFL0170I - Processing started for FASTSWITCH at 16:18:19 on 05-25-2017
.....
0173I - FASTSWITCH deletion of old primary data sets - SUCCESSFUL
.....
PFL0170I - Processing ended for FASTSWITCH at 16:18:25 on 05-25-2017
```

- Failed fast switch: PFL0172I

```
PFL0170I - Processing started for FASTSWITCH at 16:40:29 on 05-25-2017
.....
PFL0171E - UPM CONNECT FAILED for FASTSWITCH UPM RC=0008
PFL0173I - FASTSWITCH deletion of old primary data sets - BYPASSED
PFL0172E - FAILED to FASTSWITCH RC=0008 RSN=0021
PFL0170I - Processing ended for FASTSWITCH at 16:40:29 on 05-25-2017
```

CA Fast Load for DB2 for z/OS

- Functionality: Online Load Replace...Summary

- Use case: Data Replacement
 - 24/7 environment
 - Large reference tables
 - Data refresh at regular intervals
 - High data availability required
- Fast switch load: Online Load Replace
 - DATA-AVAILABLE HIGH
 - OUTPUT-CONTROL BUILD
 - RESUME NO REPLACE



For other technical insights and to consult your peers and product management, monitor our global communities:

CA Mainframe Community (<https://communities.ca.com/web/mainframe-2.0-community/welcome>)
CA Database Management Solutions for DB2 and IMS Global User Community (<https://communities.ca.com/web/ca-db2-tools-global-user-community/welcome>)

Thank You.



Manuel Gómez Burriel

DB2 Utilities Product Owner

ManuelAnibal.GomezBurriel@ca.com



in

Call for Speakers Now Open

